

# An $\mathcal{O}(m^2 \log m)$ -Competitive Algorithm for Online Machine Minimization\*

Lin Chen<sup>1</sup>, Nicole Megow<sup>2</sup>, and Kevin Schewior<sup>1</sup>

<sup>1</sup> Technische Universität Berlin, Institut für Mathematik, Berlin, Germany.  
 {lchen, schewior}@math.tu-berlin.de.

<sup>2</sup> Technische Universität München, Zentrum für Mathematik, Garching bei München, Germany.  
 nmegow@ma.tum.de

**Abstract.** We consider the online machine minimization problem in which jobs with hard deadlines arrive online over time at their release dates. The task is to determine a feasible schedule on a minimum number of machines. Our main result is a general  $\mathcal{O}(m^2 \log m)$ -competitive algorithm for the preemptive online problem, where  $m$  is the optimal number of machines used in an offline solution. This is the first improvement on an  $\mathcal{O}(\log(p_{\max}/p_{\min}))$ -competitive algorithm by Phillips et al. (STOC 1997), which was to date the best known algorithm even when  $m = 2$ . Our algorithm is  $\mathcal{O}(1)$ -competitive for any  $m$  that is bounded by a constant. To develop the algorithm, we investigate two complementary special cases of the problem, namely, laminar instances and agreeable instances, for which we provide an  $\mathcal{O}(\log m)$ -competitive and an  $\mathcal{O}(1)$ -competitive algorithm, respectively. Our  $\mathcal{O}(1)$ -competitive algorithm for agreeable instances actually produces a non-preemptive schedule, which is of its own interest as there exists a strong lower bound of  $n$ , the number of jobs, for the general non-preemptive online machine minimization problem by Saha (FSTTCS 2013), which even holds for laminar instances.

## 1 Introduction

We consider the fundamental problem of minimizing the number of machines that is necessary for feasibly scheduling jobs with release dates and hard deadlines. We consider the online variant of this problem in which every job becomes known to the online algorithm only at its release date. We denote this problem as the *online machine minimization problem*. We will show that we may restrict to the *semi-online* problem variant in which the online algorithm is given the optimal number of machines,  $m$ , in advance.

Our main result is a  $\mathcal{O}(m^2 \log m)$ -competitive algorithm for the preemptive online machine minimization problem. This is the first improvement upon a  $\mathcal{O}(\log \frac{p_{\max}}{p_{\min}})$ -competitive algorithm [17]. Our competitive ratio depends only on the optimum value,  $m$ , instead of other input parameters. In particular, it is constant if the optimum is bounded by a fixed constant. Such a result was not known even for  $m = 2$ .

**Previous results.** The preemptive semi-online machine minimization problem, in which the optimal number of machines is known in advance, has been investigated extensively by Phillips et al. [17], and there have hardly been any improvements since then. Phillips et al.

---

\* This research was supported by the German Science Foundation (DFG) under contract ME 3825/1. The third author was supported by DFG within the research training group ‘Methods for Discrete Structures’ (GRK 1408).

show a general lower bound of  $\frac{5}{4}$  and leave a huge gap to the upper bound  $\mathcal{O}(\log \frac{p_{\max}}{p_{\min}})$  on the competitive ratio for the so-called *Least Laxity First* (LLF) Algorithm. Not so surprisingly, they also rule out that the *Earliest Deadline First* (EDF) Algorithm may improve on the performance of LLF; indeed they show a lower bound of  $\Omega(\frac{p_{\max}}{p_{\min}})$ . It is a wide open question if preemptive semi-online machine minimization admits a constant competitive ratio or even a competitive ratio independent of the number of jobs or processing times, e.g., an  $f(m)$ -competitive algorithm for some function  $f$ . It is not even known whether a constant competitive algorithm exists for  $m = 2$ .

The non-preemptive problem is considerably harder than the preemptive problem. If the set of jobs arrives online over time, then no algorithm can achieve a competitive ratio sublinear in the number of jobs [18]. However, relevant special cases admit online algorithms with small constant worst-case guarantees. The problem with unit processing times was studied in a series of papers [9, 14, 15, 18, 19] and implicitly in the context of energy minimization in [4]. It has been shown that an optimal online algorithm has the exact competitive ratio  $e \approx 2.72$  [4, 9]. For non-preemptive scheduling of jobs with equal deadlines, an upper bound of 16 is given in [9]. We are not aware of any previous work on online machine minimization restricted to agreeable instances. However, in other contexts, e.g., online buffer management [13] and scheduling with power management [1, 3], it has been studied as an important and relevant class of instances.

In a closely related problem variant, an online algorithm is given extra speed to the given number of machines instead of additional unit-speed machines. The goal is to find an algorithm that requires the minimum extra speed. This problem seems much better understood and speedup factors below 2 are known (see [2, 16, 17]). However, the power of speed is much stronger than that of additional machines since it can be viewed to allow parallel processing of jobs to some extent. None of the algorithms that are known to perform well for the speed-problem, e.g., EDF and LLF, are  $f(m)$ -competitive for any function  $f$  for the machine minimization problem.

We also mention that the offline problem, in which all jobs are known in advance, can be solved optimally in polynomial time if job preemption is allowed [11]. Again, the problem complexity increases drastically if preemption is not allowed. In fact, the problem of deciding whether one machine suffices to schedule all the jobs non-preemptively is strongly NP-complete [10]. It is even open if a constant-factor approximation exists; a lower bound of  $2 - \varepsilon$  was given in [8]. The currently best known non-preemptive algorithm has an approximation factor of  $\mathcal{O}(\sqrt{(\log n)/(\log \log n)})$  [7]. Small constant factors were obtained for special cases [8, 20]. However, when slightly increasing the speed of the machines, then also the general problem can be approximated within a factor 2 [12].

**Our Contribution.** Our main contribution is a new preemptive online algorithm with a competitive ratio  $\mathcal{O}(m^2 \log m)$ , where  $m$  is the optimum number of machines. It is the first improvement upon the  $\mathcal{O}(\log \frac{p_{\max}}{p_{\min}})$ -competitive algorithm by Phillips et al. [17]. Specifically, if the optimum value  $m$  is bounded by a constant, our algorithm is  $\mathcal{O}(1)$ -competitive.

We achieve this algorithm by firstly studying two complementary special cases of the problem, namely, laminar instances and agreeable instances. In a laminar instance, if the feasible time intervals for processing (between release date and deadline) of any two jobs overlap, then one interval is completely contained in the other. In an agreeable instance, however, the relative order of release dates coincides with that of the corresponding deadline. We provide an  $\mathcal{O}(\log m)$ -competitive algorithm for laminar instances and an  $\mathcal{O}(1)$ -

competitive algorithm for agreeable instances. Then we combine both techniques to derive an  $\mathcal{O}(m^2 \log m)$ -competitive algorithm for the general problem.

It is not difficult to see that we may assume that the optimum number of machines  $m$  is known, i.e, the semi-online model, and that jobs with a small processing time relative to the entire time window (“loose” jobs) are easy to schedule. For agreeable instances we show that when scheduling “tight” jobs simply in the middle of their time windows, then there are at most  $\mathcal{O}(m)$  jobs running at the same time. Our  $\mathcal{O}(1)$ -competitive algorithm actually produces a non-preemptive schedule. This result is of its own interest as there exists a strong lower bound of  $n$  for the general non-preemptive (semi)-online machine minimization problem [18], which even holds for the special case of laminar instances.

The most difficult special case seems to be the scheduling of “tight” laminar jobs. Here we separate the job-to-machine assignment from the scheduling procedure. Upon arrival, we assign a job irrevocably to a single machine and simply run EDF on each machine individually. Thus we restrict ourselves to non-migratory schedules, which allows us to keep a better control over the remaining processing capacity in a time interval. To assign a job to a machine we consider on each machine the previously assigned smallest job that is relevant at this point in time, and we find one whose laxity is large enough to fully cover the new job’s entire time window. We then assign the new job to the same machine. Fitting a job’s *full* time window (instead of only the processing volume) into the laxity of a previously assigned job may seem rather restrictive. But jobs are “tight”, i.e., they contribute a significant processing volume to the interval, and we gain much more structure for the analysis. A greedy assignment turns out to fill up laxities too aggressively. To slow down this process we employ a more sophisticated balancing scheme in which the laxity of a job is divided into evenly-sized bins and jobs are distributed carefully over these bins. Our analysis shows that the algorithm is  $\mathcal{O}(\log m)$ -competitive.

For the general setting it is tempting to use the previous algorithms as a black box. However, it is unclear if an online partitioning into a small number of agreeable and laminar sub-instances is possible. Instead, we propose a more sophisticated variant of the assignment procedure for the laminar case, losing an additional factor of  $m$ . Instead of assigning a new job to a single machine, we form  $\mathcal{O}(m^2 \log m)$  groups of jobs. Each of these groups is then scheduled on  $\mathcal{O}(m)$  machines by generalizing the idea from the agreeable setting.

**Outline.** In Section 2, we define the problem and give first structural insights to the problem. We give an  $\mathcal{O}(1)$ -competitive algorithm for agreeable instances in Section 3 and an  $\mathcal{O}(\log m)$ -competitive algorithm for laminar instances in Section 4. In Section 5, we show how to extend the techniques for the special cases to an  $\mathcal{O}(m^2 \log m)$ -competitive algorithm for the general problem.

## 2 Preliminaries

**Problem Definition.** Given is a set of jobs  $J = \{1, 2, \dots, n\}$  where each job  $j \in J$  has a processing time  $p_j \in \mathbb{N}$ , a release date  $r_j \in \mathbb{N}$  which is the earliest possible time at which the job can be processed, and a deadline  $d_j \in \mathbb{N}$  by which it must be completed. The task is to open a minimum number of machines such that there is a feasible schedule in which no job misses its deadline. In a feasible schedule each job  $j \in J$  is scheduled for  $p_j$  units of time within the time window  $[r_j, d_j]$ . Each opened machine can process at most one job at any time, and no job is running on multiple machines at the same time. Unless stated differently, we allow job preemption, that is, a job can be preempted at any moment in time

and may resume processing later on the same or any other machine. When preemption is not allowed, then a job must run until completion once it has started.

To evaluate the performance of our online algorithms, we perform a *competitive analysis* (see e.g. [5]). We call an online algorithm  $A$   $c$ -competitive if  $m_A$  machines with  $m_A \leq c \cdot m$  suffice to guarantee a feasible solution for any instance that admits a feasible schedule on  $m$  machines.

**Notation.** For any job  $j \in J$ , the *laxity* is defined as  $\ell_j = d_j - r_j - p_j$ . We call a job  $\alpha$ -loose, for some  $\alpha < 1$ , if  $p_j \leq \alpha(d_j - r_j)$  and  $\alpha$ -tight otherwise. The (processing) interval of  $j$  is  $I(j) = [r_j, d_j]$ . For a set of jobs  $S$ , we define  $I(S) = \cup_{j \in S} I(j)$ . For  $I = \cup_{i=1}^k [a_i, b_i]$  where  $[a_1, b_1), \dots, [a_k, b_k)$  are pairwise disjoint, we define the *length* of  $I$  to be  $|I| = \sum_{i=1}^k (b_i - a_i)$ .

**Characterization of the Optimum.** For  $I$  as above, the *contribution* of a job  $j$  to  $I$  is  $C(j, I) := \max\{0, |I \cap I(j)| - \ell_j\}$ , that is, the minimum processing time that  $j$  must receive in  $I$  in any feasible schedule. The contribution of a job set  $S$  to  $I$  is the sum of the individual contributions of jobs in  $S$ , and we denote it by  $C(S, I)$ . Clearly, if  $S$  admits a feasible schedule on  $m$  machines,  $C(S, I)/|I|$  must not exceed  $m$ . Interestingly, this bound is tight, which we prove in the appendix.

**Theorem 1.** *Let  $m$  be the minimum number of machines needed to schedule instance  $J$  feasibly. Then there exists a union of intervals  $I$  with  $\lceil C(J, I)/|I| \rceil = m$  but none with  $\lceil C(J, I)/|I| \rceil > m$ .*

**Reduction to the Semi-Online Problem.** We show that we may assume that the optimum number of machines  $m$  is known in advance by losing at most a factor 4 in the competitive ratio. To do so, we employ the general idea of *doubling* an unknown parameter [6]. More specifically, we open additional machines whenever the optimum solution has doubled.

Let  $A_\rho(m)$  denote a  $\rho$ -competitive algorithm for the semi-online machine minimization problem given the optimum number of machines  $m$ . Further, denote by  $m(t)$  the minimum number of machines needed to feasibly schedule all jobs released up to time  $t$ . Then our algorithm for the online problem is as follows.

**Algorithm Double:**

- Let  $t_0 = \min_{j \in J} r_j$ . For  $i = 1, 2, \dots$  let  $t_i = \min\{t \mid m(t) > 2m(t_{i-1})\}$ .
- At any time  $t_i$ ,  $i = 0, 1, \dots$ , open  $2\rho m(t_i)$  additional machines. All jobs with  $r_j \in [t_{i-1}, t_i)$  are scheduled by Algorithm  $A_\rho(2m(t_{i-1}))$  on the machines opened at time  $t_{i-1}$ .

Since the time points  $t_0, t_1, \dots$  as well as  $m(t_0), m(t_1), \dots$  can be computed online and  $A_\rho$  is assumed to be an algorithm for the semi-online problem, this procedure can be executed online. Notice that **Double** does not preempt jobs which would not have been preempted by Algorithm  $A_\rho$ .

**Theorem 2.** *Given a  $\rho$ -competitive algorithm for (non-)preemptive semi-online machine minimization, **Double** is  $4\rho$ -competitive for (non-)preemptive online machine minimization.*

In the rest of the paper we will thus be concerned with the semi-online problem.

**Scheduling Loose Jobs.** We show that, for any fixed  $\alpha < 1$ ,  $\alpha$ -loose jobs are easy to handle via a simple greedy algorithm called *Earliest Deadline First* (EDF) that schedules at any time  $m' = \rho m$  unfinished jobs with the smallest deadline. If every job is feasibly scheduled, it is  $\rho$ -competitive.

**Theorem 3.** *Let  $\alpha \in (0, 1)$ . EDF is  $1/(1 - \alpha)^2$ -competitive for any instance that consists only of  $\alpha$ -loose jobs.*

As we aim for asymptotical competitive ratios in this paper, we can from now on assume that all jobs are  $\alpha$ -tight for a fixed  $\alpha \in (0, 1)$ .

### 3 A Constant-Competitive Ratio for Agreeable Instances

Consider an instance  $J$  in which any two jobs  $j, j' \in J$  are *agreeable*, that is,  $r_j < r_{j'}$  implies  $d_j \leq d_{j'}$ . We also call  $J$  agreeable. For such instances, we derive an  $\mathcal{O}(1)$ -competitive online algorithm. By Theorem 2, we may assume that the optimum value  $m$  is known in advance. Using Theorem 3, we can choose some constant  $\alpha \in (0, 1)$  and schedule  $\alpha$ -loose jobs by EDF on a separate set of machines. It remains to show an  $\mathcal{O}(1)$ -competitive algorithm for  $\alpha$ -tight jobs.

We define the  $\beta$ -interval of a job  $j$ , for some  $\beta \leq 1/2$ , as  $I_\beta(j) := [r_j + \beta\ell_j, d_j - \beta\ell_j]$ . We use the following simple non-preemptive algorithm.

**Algorithm MediumFit:** Each job  $j$  is scheduled in the middle of its feasible time window, i.e., during  $I_{1/2}(j)$ . If for some job  $j$  there is no vacant machine during  $I_{1/2}(j)$ , then we open a new machine.

For the analysis of **MediumFit**, let two jobs  $j, j'$  be  $\beta$ -agreeable if they are agreeable and their  $\beta$ -intervals have a non-empty intersection. We give an upper bound on the number of jobs that can be  $\beta$ -agreeable with a single job. In fact, we will prove a stronger result than actually needed for this section, but it will be useful for analyzing the general algorithm in Section 5.

**Lemma 1.** *Consider a job  $j$  and some fixed  $\alpha \in (0, 1)$  and  $\beta \in (0, 1/2]$ . Then there exist at most  $\mathcal{O}(m)$   $\alpha$ -tight jobs  $j'$  with  $\ell_{j'} \leq \ell_j$  such that  $j$  and  $j'$  are  $\beta$ -agreeable.*

*Proof.* Consider all such  $j'$ . We estimate their contributions to the interval(s)

$$I = [r_j - 2\ell_j, r_j + 2\ell_j] \cup [d_j - 2\ell_j, d_j + 2\ell_j],$$

which has a total length of at most  $8\ell_j$ . There are two possibilities for  $j'$ .

- Case 1: We have  $|I(j')| \geq 2\ell_j$ . As  $I(j')$  contains  $r_j$  or  $d_j$ , we have  $|I \cap I(j')| \geq 2\ell_j$ . Given that  $\ell_{j'} \leq \ell_j$ ,  $j'$  contributes at least  $2\ell_j - \ell_{j'} \geq \ell_j$  to  $I$ .
- Case 2: We have  $|I(j')| < 2\ell_j$ . As a consequence  $I(j') \subseteq I$ . Observe that  $|I(j')| \geq \beta\ell_j$  holds because  $j$  and  $j'$  are  $\beta$ -agreeable. As  $j'$  is  $\alpha$ -tight, its contribution to  $I$  is at least  $\alpha\beta\ell_j$ .

Let  $n_1$  and  $n_2$  be the number of jobs corresponding to the above two cases, respectively. Then the contribution of the  $n_1 + n_2$  jobs to  $I$  is at least

$$(n_1 + \alpha\beta n_2) \cdot \ell_j \geq (n_1 + n_2) \cdot \alpha\beta\ell_j.$$

Using Theorem 1, the total contribution is upper bounded by  $m|I| \leq 8m\ell_j$ , implying  $n_1 + n_2 \leq 8m/\alpha\beta = \mathcal{O}(m)$ .  $\square$

Now, consider the **MediumFit** schedule. Any two jobs that are processed at the same time are  $\beta$ -agreeable for  $\beta = 1/2$ . By Lemma 1, the number of such jobs—and, thus, the number of required machines—is  $\mathcal{O}(m)$ . Therefore, we have an  $\mathcal{O}(1)$ -competitive algorithm for agreeable instances.

Notice also that our final schedule is non-preemptive. **MediumFit** (for  $\alpha$ -tight jobs) is by definition a non-preemptive algorithm. EDF (for  $\alpha$ -loose jobs) on agreeable instances never preempts jobs that have already started because the jobs released later have a larger deadline. Since we compete with a weaker optimum in the non-preemptive setting, our result carries over to this setting.

**Theorem 4.** *On agreeable instances, there is an  $\mathcal{O}(1)$ -competitive algorithm for preemptive and non-preemptive online machine minimization.*

## 4 $\mathcal{O}(\log m)$ -Competitiveness for Laminar Instances

In this section, we consider *laminar* instances in which for any two jobs  $j, j'$  with *overlapping* intervals, that is,  $|I(j) \cap I(j')| > 0$ , holds that either  $I(j) \subseteq I(j')$  or  $I(j') \subseteq I(j)$ . We prove the following result.

**Theorem 5.** *On laminar instances, there is an  $\mathcal{O}(\log m)$ -competitive algorithm for online machine minimization.*

We assume that jobs are indexed from 1 to  $n$  in accordance with their release dates, that is,  $j < j'$  implies  $r_j \leq r_{j'}$ . If  $r_j = r_{j'}$ , we assume that  $j < j'$  implies  $d_j \geq d_{j'}$ . We say that  $j$  *dominates*  $j'$  (denoted as  $j \succ j'$ ) if  $j < j'$  and  $I(j) \supseteq I(j')$ . We also say  $j'$  *is dominated by*  $j$  and denote this as  $j' \prec j$ .

### 4.1 Description of the Algorithm

By Theorems 2 and 3, it again suffices to consider semi-online scheduling of  $\alpha$ -tight jobs for some fixed  $\alpha \in (0, 1)$ .

**Job Assignment.** We open  $m'$  machines and will later show that we can choose  $m' = \mathcal{O}(m \log m)$ . At every release date, we immediately assign a new job to its machine and we never revoke this decision. Thus, we will obtain a non-migratory schedule.

The assignment procedure is as follows. We consider an unassigned new job  $j$ . If there is a machine that is “surely free” at time  $r_j$ , that is, there is no job  $j'$  whose time window  $I(j')$  contains  $r_j$ , then we assign  $j$  to this machine. Otherwise we do the following.

We consider all previously assigned jobs  $j'$  with  $r_j \in I(j')$ . Since the instance is laminar, there is on each machine a unique  $\prec$ -*minimal* job among them (unless intervals have identical time intervals, in which case we pick the one with the smallest index), that is, a job that is not dominating any other job on this machine. Essentially, this is the job that has the smallest interval length  $I(j')$  on a machine. We call each such job *candidate*. Now, consider the laminar chain  $c_1(j) \prec \dots \prec c_{m'}(j)$  of all candidates at time  $r_j$ . Again, this chain exists because the instance is laminar. We call  $c_i(j)$  the  $i$ -th *candidate* of  $j$ . Notice that  $I(j) \subseteq I(c_i(j))$  for all  $i$ , that is,  $j$  is dominated by all its candidates.

We want to assign  $j$  to one of the machines and need to ensure that it can feasibly be scheduled. To that end, we will find a candidate  $c_i(j)$  that has enough capacity (i.e., laxity) in its time window to fit the *full* time interval  $I(j)$  and assign  $j$  to this candidate

$c_i(j)$ . Clearly, we have to carefully keep track of the contribution of previously assigned jobs to  $c_i(j)$  which reduces the remaining capacity (i.e., the laxity) within the candidate's time window. However, a greedy assignment of jobs to their  $\prec$ -minimal candidates fails, and we employ a more sophisticated balancing scheme.

We partition the laxity of each candidate's time window into  $m'$  equally-sized bins. To assign  $j$ , we select the smallest  $i$  such that the  $i$ -th bin of the  $i$ 'th candidate  $c_i(j)$  still has a remaining capacity of at least  $|I(j)|$ . If there is no such  $i$  then the assignment of  $j$  *fails*.

To make this assignment more precise and to enable us to analyze the procedure later, we introduce some notation. We denote a job  $j$  that has been assigned to its  $i$ -th candidate  $c_i(j)$  as the  $i$ -th *user* of  $c_i(j)$ . We denote the set of all  $i$ -th users of a job  $j'$  by  $U_i(j')$ , for any  $i$ . All  $i$ -th users of  $j'$  will contribute to the  $i$ -th bin of  $j'$ . As stated earlier, to assign  $j$ , we select the smallest  $i$  such that the  $i$ -th bin of  $c_i(j)$  still has a remaining capacity of at least  $|I(j)|$ , that is,

$$|I(U_i(c_i(j))) \cup I(j)| \leq \frac{\ell_{c_i(j)}}{m'}. \quad (1)$$

If we find such an  $i$  and assign  $j$ , then the capacity of the  $i$ -th bin of  $c_i(j)$  reduces by  $|I(j)|$ . Otherwise, the assignment of  $j$  *fails*.

**Scheduling.** Given a job-to-machine assignment, we run EDF on each machine separately.

## 4.2 Analysis of the Algorithm

We first show that our algorithm obtains a feasible schedule if no job assignment has failed. Then we give a proof of the fact that the job assignment never fails on instances that admit a feasible schedule.

**Lemma 2.** *If the job assignment does not fail, our algorithm schedules feasibly.*

*Proof.* Consider some job  $j$  and note that, whenever it is preempted at some time  $t$ , there must be a user  $j'$  of  $j$  with  $t \in I(j')$ . According to Inequality (1),  $j$  is thus preempted for no longer than

$$\sum_{i=1}^{m'} |I(U_i(j))| \leq m' \cdot \frac{\ell_j}{m'} = \ell_j.$$

Therefore, it can be processed for  $p_j$  time units.  $\square$

It remains to show that the job assignment never fails for some  $m' = \mathcal{O}(m \log m)$ . The proof idea is as follows: We assume that the algorithm fails to assign a job. We select a critical job set, the *failure set*, which is the set of jobs that contribute (directly or indirectly) to the  $m'$ -th bin of the  $m'$ -th candidate of the failing job. We take into account the direct contribution by users of that bin and the indirect contribution by jobs that filled lower-indexed bins of lower-indexed candidates and thus indirectly contributed to the filling of the last possible bin. With this failure set we derive a contradiction using a load argument.

Let  $j^*$  be a job whose assignment fails. We iteratively define its *failure set*  $F \subseteq J$  as the union of the  $m' + 1$  different sets  $F_0, \dots, F_{m'}$ . To define the base set  $F_0$ , we also use auxiliary sets  $F_0^0, \dots, F_0^{m'}$ .

We initialize  $F_0^{m'} := \{j^*\}$ . Given  $F_0^i$ , we construct  $F_i$  and  $F_0^{i-1}$  in the following way. First,  $F_i$  is defined to be the set of all  $\prec$ -maximal  $i$ -th candidates of jobs in  $F_0^i$ . Subsequently,

$F_0^{i-1}$  is constructed by incrementing the set of users in  $F_0^i$  by the  $i$ 'th users of jobs in  $F_i$ . Formally,

$$F_i := M_{\prec}(\{c_i(j) \mid j \in F_0^i\})$$

$$\text{and } F_0^{i-1} := F_0^i \cup \bigcup_{j \in F_i} U_i(j),$$

where  $M_{\prec}$  is the operator that chooses the  $\prec$ -maximal elements from a set of jobs:  $M_{\prec}(S) := \{j \in S \mid \nexists j' : j \prec j'\}$ . After  $m'$  such iterations, that is, when  $F_0^0, F_1, \dots, F_{m'}$  have been computed, we set  $F_0 := M_{\prec}(F_0^0)$ . Our failure set is  $F_0 \cup F_1 \cup \dots \cup F_{m'}$ .

We show some properties of failure sets which will be useful for the proofs later. To that end, we define the following notation. For two job sets  $S_1$  and  $S_2$ , we write  $S_1 \prec S_2$  if, for each  $j_1 \in S_1$ , there is a  $j_2 \in S_2$  with  $j_1 \prec j_2$ . Moreover, a job set  $S$  is said to  $\gamma$ -block a job  $j$  if there is a subset  $S' \subseteq S$  with  $S' \prec \{j\}$  and  $|I(S')| \geq \gamma \ell_j$ . Finally,  $S_1$   $\gamma$ -blocks  $S_2$  if it  $\gamma$ -blocks every job in  $S_2$ .

**Lemma 3.** *Every failure set  $F$  has the following properties.*

- (i) *We have  $F_0 \prec \dots \prec F_{m'}$ .*
- (ii) *The sets  $F_0, \dots, F_{m'}$  are pairwise disjoint.*
- (iii) *For all  $i = 1, \dots, m'$ ,  $F_0 \frac{1}{m'}$ -blocks  $F_i$ .*

*Proof.* To see (i), we define  $C_i := \{c_i(j) \mid j \in F_0^i\}$ , for every  $1 \leq i \leq m'$ , and  $C_0 := F_0^0$ . We first show  $C_{i-1} \prec C_i$  for every  $1 \leq i \leq m'$ . For  $i = 1$ , this directly follows from the construction. Consider  $2 \leq i \leq m'$ . Let  $j$  be an arbitrary job in  $C_{i-1}$ , which is then an  $(i-1)$ -th candidate of some job in  $F_0^{i-1}$ , say, job  $j'$ . According to the construction, we have  $j' \in F_0^i$ , or  $j'$  is the  $i$ -th user of some job in  $C_i$ . In both cases,  $C_i$  contains the  $i$ -th candidate of job  $j'$ , which dominates job  $j$ . Hence,  $C_{i-1} \prec C_i$ . Since  $F_{i-1}$  and  $F_i$  are obtained from  $C_{i-1}$  and  $C_i$  only by deleting dominated jobs,  $F_{i-1} \prec F_i$  follows.

Consider property (ii) and suppose there is some  $j \in F_i \cap F_{i'}$  for some  $i < i'$ . Then, by (i), there is also a job  $j' \in F_{i'}$  with  $j \prec j'$ , contradicting the fact that  $F_{i'}$  only contains  $\prec$ -maximal elements.

Consider property (iii). Let  $1 \leq i \leq m'$  and  $j' \in F_i$  be an arbitrary job. We show that there is some  $F'_0 \subseteq F_0$  with  $F'_0 \prec \{j'\}$  and  $|I(F'_0)| \geq \ell_{j'}/m'$ . By the construction of  $F_i$ , we have  $j' = c_i(j)$  for some  $j \in F_0^i$ . Recall the construction of  $F_0^i$ : We have  $j = j^*$ , or  $j$  is the  $i'$ -th user of some job where  $i' \geq i + 1$ . In both cases, as our algorithm assigns  $j$  to the machine of the smallest possible candidate,  $j$  could not be assigned to the same machine as  $j'$ , which is its  $i$ -th candidate. Thus, when our algorithm assigned job  $j$ , the capacity of the  $i$ -th bin of  $j'$  did not have enough remaining capacity for  $j$  to become its user. Hence,

$$|I(U_i(j') \cup \{j\})| > \ell_{j'}/m'.$$

Note that  $U_i(j') \cup \{j\} \subseteq F_0^0$  holds and that  $U_i(j') \cup \{j\}$  consists of pairwise disjoint jobs. As  $F_0$  is obtained by only deleting jobs in  $F_0^0$  that are dominated by others, we can always find  $F'_0 \subseteq F_0$  with  $|I(F'_0)| \geq \ell_{j'}/m'$ .  $\square$

Using a load argument, we can even derive stronger blocking relations.

**Lemma 4.** *For any fixed  $\alpha \in (0, 1)$ , there exists  $k = \mathcal{O}(m)$  with the following property. If  $F_i$   $\gamma$ -blocks  $F_h$ , then  $F_{i+k}$   $2\gamma$ -blocks  $F_h$ , for all  $h > i + k$ .*



*Proof.* Consider some arbitrary  $j \in F_h$ . For every  $i' \in \{i, \dots, i+k\}$ , we denote by  $F_{i'}^*$  the subset of  $F_{i'}$  consisting of jobs that are dominated by  $j$ . We prove that, for  $k = \mathcal{O}(m)$  the inequality  $|I(F_{i+k}^*)| \geq 2|I(F_i^*)|$  holds. By Lemma 3 (i) and (ii), we have the chain  $I(F_i^*) \subseteq \dots \subseteq I(F_{i+k}^*)$ . As all of these jobs are  $\alpha$ -tight, we can thus lower bound the processing-time contribution of  $F_{i'}^*$  to  $I(F_{i+k}^*)$  by  $\alpha|I(F_{i'}^*)| \geq \alpha|I(F_i^*)|$ , for all  $i' \in \{i, \dots, i+k\}$ . The contributions of all these jobs sum up to at least  $k\alpha|I(F_i^*)|$ , which must not exceed  $m|I(F_{i+k}^*)|$  by Theorem 1. Hence, for  $k > 2m/\alpha = \mathcal{O}(m)$  we have  $|I(F_{i+k}^*)| \geq 2|I(F_i^*)|$ , which implies the claim.  $\square$

**Lemma 5.** *There exists  $m' = \mathcal{O}(m \log m)$  such that the assignment never fails.*

*Proof.* We will choose  $m' = \mathcal{O}(m \log m)$  such that we get a contradiction to the existence of the failure set  $F$ . Using Lemma 3 (iii) and Lemma 4, we can choose  $\lambda = \mathcal{O}(m \log m')$  such that  $F_\lambda$  2-blocks every  $F_i$  with  $i > \lambda$ . Now we lower bound the total contribution of these  $F_i$  to  $I(F_\lambda)$ : For any  $j \in F_i$  where  $i \geq \lambda$ , its contribution to  $I(F_\lambda)$  is

$$|I(F_\lambda) \cap I(j)| - \ell_j \geq |I(F_\lambda) \cap I(j)|/2$$

since  $|I(F_\lambda) \cap I(j)| \geq 2\ell_j$ . Thus we get as total contribution of  $F_i$  to  $F_\lambda$  at least

$$\sum_{j \in F_i} \frac{|I(F_\lambda) \cap I(j)|}{2} \geq \frac{1}{2} \cdot \left| I(F_\lambda) \cap \left( \bigcup_{j \in F_i} I(j) \right) \right| = \frac{|I(F_\lambda) \cap I(F_i)|}{2} \geq \frac{|I(F_\lambda)|}{2},$$

where the last inequality follows from  $F_\lambda \prec F_i$ . Therefore, the total contribution of all  $F_i$  for  $i > \lambda$  is at least  $(m' - \lambda)|I(F_\lambda)|/2$ , which is at most  $m|I(F_\lambda)|$  by Theorem 1. Thus, we must have  $m' \leq \lambda + 2m = \mathcal{O}(m \log m')$ , and appropriately choosing  $m' = \mathcal{O}(m \log m)$  yields a contradiction.  $\square$

As stated earlier, combining this lemma with Lemma 2 shows Theorem 5.

## 5 $\mathcal{O}(m^2 \log m)$ -Competitiveness in the General Case

Now, we generalize the methods introduced above and prove our main result.

**Theorem 6.** *There is an  $\mathcal{O}(m^2 \log m)$ -competitive algorithm for online machine minimization.*

It is not clear if an online separation into a reasonable number of laminar and agreeable sub-instances exists such that we could utilize the previous algorithms as black boxes. Nevertheless, we take over the notion of domination from the laminar case and slightly extend it using the previously introduced definition of  $\delta$ -intervals  $I_\delta(j) := [r_j + \delta\ell_j, d_j - \delta\ell_j]$ . We say job  $j$   $\delta$ -dominates job  $j'$  if job  $j$  dominates  $j'$  and we have  $I(j') \cap I_\delta(j) \neq \emptyset$ . The presented algorithm works for any fixed  $\delta \in (0, 1/2)$ .

### 5.1 Description of the Algorithm

Again, by Theorems 2 and 3, it suffices to consider semi-online scheduling of  $\alpha$ -tight jobs where  $\alpha$  is fixed. We describe an  $\mathcal{O}(m^2 \log m)$ -competitive algorithm. In its structure, our algorithm mainly resembles the previous algorithm for laminar instances and it incorporates ideas from the one for agreeable instances.

**Job Assignment.** Instead of assigning each job immediately to a fixed machines, we form  $g$  initially empty *groups* of jobs. We will later choose  $g = \mathcal{O}(m^2 \log m)$  and show that each of these groups can be scheduled on  $\mathcal{O}(m)$  machines. Upon its release, each job  $j$  is immediately assigned to one of the groups based on the assignment of the previously released  $\delta$ -dominating jobs. If there exists a group in which no job  $\delta$ -dominates  $j$ , then  $j$  is assigned to such a group. Otherwise, in contrast to the laminar case, there might be multiple  $\prec$ -minimal such jobs in each group. Among these jobs we now pick the one with the *minimum laxity* from each group, and we denote this set of  $g$  jobs as  $D$ .

Like in the laminar case, we choose  $\tau$  *candidates*, where  $\tau$  is independent of  $j$  and will be specified later. Again, these jobs will form a chain  $c_1(j) \prec \dots \prec c_\tau(j)$ , where we call  $c_i(j)$  the  $i$ -th candidate of  $j$ , for all  $i$ . Recall that this step is straightforward for a laminar instance as the jobs in  $D$  form such a chain already. Here, however, this chain is constructed in an iterative way, starting with iteration 0: In iteration  $i$ , we define  $c_{\tau-i}(j)$  to be a maximum-laxity job among those jobs in  $D$  dominated by all jobs selected so far. If no such jobs exist any more before we have found the  $\tau$ -th candidate, the job assignment fails.

Analogous to the laminar case, we pick a candidate  $c_i(j)$  and assign  $j$  to the same group. As before,  $j$  will be called an  $i$ -th *user* of  $c_i(j)$  after this assignment, and the set of  $h$ -th users of some job  $j'$  is denoted by  $U_h(j')$ , for all  $h$ . We open  $\tau$  equally-sized bins for each job, where we again pack intervals of its  $h$ -th users into the  $h$ -th bin, for all  $h$ . In contrast to the laminar case, the sum of the bin sizes of a job is only a small fraction of its laxity: Here, we choose the smallest candidate  $c_i(j)$  such that

$$|I(U_i(c_i(j))) \cup I(j)| \leq \frac{\ell_{c_i(j)}}{qm\tau},$$

where  $q$  is a constant yet to be determined. The job assignment fails if we do not find such a candidate. Note that, again in contrast to the laminar case, users of the same job may overlap, that is,  $|I(U_i(c_i(j))) \cup I(j)|$  is only a lower bound on the sum of the individual interval lengths of jobs in  $U_i(c_i(j)) \cup \{j\}$ .

**Scheduling.** Consider a group and let  $S \subseteq J$  be the jobs assigned to it. We schedule  $S$  on  $\mathcal{O}(m)$  machines in the following way. Similarly to our algorithm for agreeable instances, each job  $j \in S$  is scheduled only within  $I_\delta(j)$ . More precisely, at any time  $t$  consider all the unfinished jobs  $j \in S$  with  $t \in I_\delta(j)$ . Among all these jobs, we schedule exactly the  $\prec$ -minimal jobs, each one on a separate machine. Note that our algorithm for laminar instances also schedules exactly the  $\prec$ -minimal jobs on each machine, which coincides with EDF for laminar instances.

## 5.2 Analysis of the Algorithm

As in the laminar case, we first show that, if the job assignment never fails, we obtain a feasible schedule for all jobs. We need the following auxiliary lemma.

**Lemma 6.** *Let  $j$  and  $j'$  be two jobs that do not dominate one another but both  $\delta$ -dominate some job  $j^*$ . Then  $j$  and  $j'$  are  $(\delta/2)$ -agreeable.*

*Proof.* We show that

$$t^* := \frac{r_{j^*} + d_{j^*}}{2} \in I_{\frac{\delta}{2}}(j). \quad (2)$$

If  $t^* \in I_\delta(j)$  holds, Equation (2) directly follows. Otherwise, let  $I$  be the largest interval that contains  $t^*$  and is a subset of  $I(j^*) \setminus I_\delta(j)$ . Note that either  $r_{j^*}$  borders  $I$  to the left or

$d_{j^*}$  borders  $I$  to the right. W.l.o.g. assume that  $I = [t, d_{j^*})$  for some  $t > r_{j^*}$ . Now we have  $t^* - t < |I|/2 \leq \delta \ell_j/2$ . As  $I_{\delta/2}(j)$  is obtained from  $I_\delta(j)$  by enlarging it by  $\delta \ell_j/2$  in both directions, Equation (2) follows.

By the same argument, we can prove  $t^* \in I_{\delta/2}(j')$ , and the claim follows.  $\square$

Now we can prove the following statement, where we will pick a suitable  $q$ .

**Lemma 7.** *If the job assignment does not fail, our algorithm produces a feasible schedule on  $\mathcal{O}(mg)$  machines.*

*Proof.* Consider an arbitrary job  $j$ , assigned to group  $h$ . We first prove that  $j$  receives a processing time of  $p_j$ , that is, it is preempted for at most  $(1 - 2\delta)\ell_j$  within  $I_\delta(j)$ . Consider an arbitrary time  $t \in I_\delta(j)$  when  $j$  is preempted, and let  $S_t$  be those jobs  $j'$  from group  $h$  with  $t \in I_\delta(j')$ . According to the scheduling rule,  $j$  is not a  $\prec$ -minimal job in  $S_t$ , so there exist jobs in  $S_t$  which are dominated, and thus  $\delta$ -dominated, by  $j$ . Among them, let  $j^*$  be a  $\prec$ -maximal one, and consider the time when job  $j^*$  is assigned to group  $h$ . Then job  $j$  is one of the  $\prec$ -minimal jobs that  $\delta$ -dominate  $j^*$  in group  $h$ . According to the algorithm, among all the  $\prec$ -minimal jobs that  $\delta$ -dominate  $j^*$  in group  $h$ , the one of the least laxity is picked and added to the chain of  $\tau$  candidates of  $j^*$ , and  $j^*$  is eventually assigned (as a user) to this candidate. Thus, job  $j^*$  is either a user of job  $j$ , or a user of some other job  $j'$  with  $\ell_{j'} \leq \ell_j$  which is also a  $\prec$ -minimal job that  $\delta$ -dominates  $j^*$ .

With the above argument, we conclude that job  $j$  is preempted either due to one of its own users or a user of some other job  $j'$  from group  $h$  such that  $\ell_{j'} \leq \ell_j$ . By our job assignment, this preempts  $j$  by at most  $\ell_j/(qm)$  or  $\ell_{j'}/(qm) \leq \ell_j/(qm)$ , respectively. We claim that there are at most  $\mathcal{O}(m)$  such  $j'$ . To see why, notice that according to Lemma 6 jobs  $j$  and  $j'$  are  $\delta/2$ -agreeable. As  $\ell_{j'} \leq \ell_j$ , according to Lemma 1 the claim follows. Thus in total job  $j$  is preempted by at most  $\mathcal{O}(m) \cdot \ell_j/(qm) \leq (1 - 2\delta)\ell_j$ , with  $q$  being a sufficiently large constant.

It remains to show that, in the same group  $h$ , we process at most  $\mathcal{O}(m)$  jobs at any time  $t$ . This is straightforward: Let  $S_t$  be as above. Note that we always process the  $\prec$ -minimal jobs in  $S_t$ , and any two of them are  $\delta$ -agreeable. Hence, according to Lemma 1, there are at most  $\mathcal{O}(m)$  such jobs.  $\square$

We now show that the job assignment never fails. In contrast to the laminar case, we also have to make sure that it never fails at finding a chain of candidates.

**Lemma 8.** *We can choose  $\tau = \Omega(g/m)$ , and our algorithm will never fail at finding a chain of  $\tau$  candidates.*

*Proof.* Let  $S$  be some set of jobs that all  $\delta$ -dominate the same job  $j^*$ . Let job  $j$  be the one with maximum laxity in  $S$ . Note that it suffices to show that, except for  $\mathcal{O}(m)$  jobs, all the jobs in  $S$  are dominated by  $j$ . Then, every time when we pick a new candidate, we lose  $\mathcal{O}(m)$  jobs, and with some  $\tau = \Omega(g/m)$  the assignment procedure does not fail when selecting the chain of candidates.

Consider any job  $j' \in S$  which is not dominated by  $j$ . If job  $j'$  does not dominate  $j$  either, then due to the fact that jobs  $j$  and  $j'$  both  $\delta$ -dominate  $j^*$ ,  $j'$  is  $(\delta/2)$ -agreeable with  $j$  according to Lemma 6. By applying Lemma 1, we get that there are at most  $\mathcal{O}(m)$  such jobs  $j'$ . Otherwise  $j'$  dominates  $j$ . As we have  $\ell_{j'} \leq \ell_j$  and  $I(j') \supseteq I(j)$  for each such  $j'$ , it follows that  $j'$  contributes at least  $p_j$  to  $I(j)$ . Using that every job is  $\alpha$ -tight, Theorem 1 now allows us to bound the number of these jobs by  $m/\alpha = \mathcal{O}(m)$ , too.  $\square$

The remainder of the proof again relies on selecting a failure set whenever a job assignment fails and deriving a contradiction from that. The arguments are the same as in Section 4. Thus, we omit them here and refer to the appendix.

## References

1. Albers, S., Müller, F., Schmelzer, S.: Speed scaling on parallel processors. *Algorithmica* 68(2), 404–425 (2014)
2. Anand, S., Garg, N., Megow, N.: Meeting deadlines: How much speed suffices? In: *Proc. of ICALP*. pp. 232–243 (2011)
3. Angel, E., Bampis, E., Chau, V.: Low complexity scheduling algorithms minimizing the energy for tasks with agreeable deadlines. *Discrete Appl. Math.* 175, 1–10 (2014)
4. Bansal, N., Kimbrel, T., Pruhs, K.: Speed scaling to manage energy and temperature. *J. ACM* 54(1) (2007)
5. Borodin, A., El-Yaniv, R.: *Online Computation and Competitive Analysis*. Cambridge University Press (1998)
6. Chrobak, M., Kenyon-Mathieu, C.: SIGACT news online algorithms column 10: Competitive-ness via doubling. *SIGACT News* 37(4), 115–126 (2006)
7. Chuzhoy, J., Guha, S., Khanna, S., Naor, J.: Machine minimization for scheduling jobs with interval constraints. In: *Proc. of FOCS*. pp. 81–90 (2004)
8. Cieliebak, M., Erlebach, T., Hennecke, F., Weber, B., Widmayer, P.: Scheduling with release times and deadlines on a minimum number of machines. In: *Proc. of TCS*. pp. 217–230 (2004)
9. Devanur, N.R., Makarychev, K., Panigrahi, D., Yaroslavtsev, G.: Online algorithms for machine minimization. *CoRR* abs/1403.0486 (2014)
10. Garey, M.R., Johnson, D.S.: Two-processor scheduling with start-times and deadlines. *SIAM J. Comput.* 6(3), 416–426 (1977)
11. Horn, W.A.: Some simple scheduling algorithms. *Naval Research Logistics Quarterly* 21(1), 177–185 (1974)
12. Im, S., Li, S., Moseley, B., Torng, E.: A dynamic programming framework for non-preemptive scheduling problems on multiple machines. In: *Proc. of SODA*. pp. 1070–1086 (2015)
13. Jež, L., Li, F., Sethuraman, J., Stein, C.: Online scheduling of packets with agreeable deadlines. *ACM Trans. Algorithms* 9(1), 5:1–5:11 (2012)
14. Kao, M.J., Chen, J.J., Rutter, I., Wagner, D.: Competitive design and analysis for machine-minimizing job scheduling problem. In: *Proc. of ISAAC*. pp. 75–84 (2012)
15. Kleywegt, A.J., Nori, V.S., Savelsbergh, M.W.P., Tovey, C.A.: Online resource minimization. In: *Proc. of SODA*. pp. 576–585 (1999)
16. Lam, T.W., To, K.K.: Trade-offs between speed and processor in hard-deadline scheduling. In: *Proc. of SODA*. pp. 623–632 (1999)
17. Phillips, C.A., Stein, C., Torng, E., Wein, J.: Optimal time-critical scheduling via resource augmentation. *Algorithmica* 32(2), 163–200 (2002)
18. Saha, B.: Renting a cloud. In: *Proc. of FSTTCS*. pp. 437–448 (2013)
19. Shi, Y., Ye, D.: Online bin packing with arbitrary release times. *Theor. Comput. Sci.* 390(1), 110–119 (2008)
20. Yu, G., Zhang, G.: Scheduling with a minimum number of machines. *Oper. Res. Lett.* 37(2), 97–101 (2009)

## A Missing Material from Section 2

### A.1 Proof of Theorem 1

*Proof (Theorem 1).* It is easy to see that  $C(J, I) \leq m|I|$  since the processing volume of  $C(J, I)$  must be finished within  $I$  of length  $|I|$  in any feasible solution, while the optimal solution can finish at most  $m|I|$  by using  $m$  machines.

It remains to show that there exists some  $I^*$  such that  $C(J, I^*)/|I^*| > m - 1$ . Notice that the offline scheduling problem can be formulated as a flow problem, and there exists an optimal schedule which makes decisions only at integral points in time. Consider any such schedule that uses  $m$  machines. We denote by  $\chi_h$  the number of time slots  $[t, t + 1)$  during which exactly  $h$  machines are occupied. Thus, for any optimal schedule on  $m$  machines that makes decisions only at integral points in time, we obtain a vector  $\chi = (\chi_m, \dots, \chi_0)$ . We define a lexicographical order  $<$  on these vectors: we have  $\chi < \chi'$  if and only if there exists some  $h$  with  $\chi_h < \chi'_h$  and  $\chi_i = \chi'_i$ , for all  $i < h$ . We pick the optimal schedule whose corresponding vector is the smallest with respect to  $<$ .

We now construct a directed graph  $G = (V, A)$  based on the schedule we have picked. Let  $V = \{v_0, \dots, v_{d_{\max}-1}\}$  where  $v_t$  represents the slot  $[t, t + 1)$ . Let  $\phi_t$  be the number of machines occupied during  $[t, t + 1)$ . An arc from  $v_i$  to  $v_k$  exists iff  $\phi_i \geq \phi_k$  and there exists some job  $j$  with  $[i, i + 1) \cup [k, k + 1) \subseteq I(j)$  whereas  $j$  is processed in  $[i, i + 1)$  but not in  $[k, k + 1)$ . Intuitively, an arc from  $v_i$  to  $v_k$  implies that one unit of the workload in  $[i, i + 1)$  could be carried to  $[k, k + 1)$ .

We claim that in  $G$  there is no (directed) path which starts from some  $v_i$  with  $\phi_i = m$ , and ends at some  $v_\ell$  with  $\phi_\ell \leq m - 2$ . Suppose there exists such a path, say  $(v_{i_1}, \dots, v_{i_\ell})$  with  $\phi_{i_1} = m$  and  $\phi_{i_\ell} \leq m - 2$ . Then we alter the schedule such that we move one unit of the workload from  $[i_s, i_s + 1)$  to  $[i_{s+1}, i_{s+1} + 1)$ , for all  $s < \ell$ . By doing so,  $\phi_{i_1}$  decreased and  $\phi_{i_\ell}$  increased by 1 each. By  $\phi_{i_1} = m$  and  $\phi_{i_\ell} \leq m - 2$ , we get that  $\chi_m$  decreases by 1, contradicting the choice of the schedule.

Consider  $V_m = \{v_i \mid \phi_i = m\}$  and let  $P(V_m) = \{v_{i_1}, \dots, v_{i_\ell}\}$  be the set of vertices reachable from  $V_m$  via a directed path (trivially,  $V_m \subseteq P(V_m)$ ). The above arguments imply that for any  $v_i \in P(V_m)$ , it holds that  $\phi_i \geq m - 1$ . We claim that  $\sum_j C(j, I^*) = \sum_h \phi_{i_h}$  for  $I^* = \{[i_h, i_h + 1) \mid 1 \leq h \leq \ell\}$ . Note there is no arc going out from  $P(V_m)$  (otherwise the endpoint of the arc is also reachable and would have been included in  $P(V_m)$ ). That is, we cannot move out any workload from  $I^*$ , meaning that the contribution of all jobs to  $I^*$  is included in the current workload in  $I^*$ , i.e.,  $\sum_h \phi_{i_h}$ . Thus,

$$\frac{C(J, I^*)}{|I^*|} \geq \frac{\sum_h \phi_{i_h}}{|I^*|} = \frac{\sum_h \phi_{i_h}}{|P(V_m)|}.$$

Notice that  $\phi_{i_h}$  is either  $m$  or  $m - 1$ , and among them there is at least one of value  $m$ , thus the right-hand side is strictly larger than  $m - 1$ , i.e.,  $C(J, I^*)/|I^*| > m - 1$ .  $\square$

We remark that this result can be proven also using LP duality.

### A.2 Proof of Theorem 2

*Proof (Theorem 2).* Let  $t_0, t_1, \dots, t_k$  denote the times at which Double opens new machines. Double schedules the jobs  $J_i = \{j \mid r_j \in [t_i, t_{i+1})\}$ , with  $i = 0, 1, \dots, k$ , using Algorithm  $A_\rho(2m(t_i))$  on  $2\rho m(t_i)$  machines exclusively. This yields a feasible schedule since

an optimal solution for  $J_i$  requires at most  $m(t_{i+1} - 1) \leq 2m(t_i)$  machines and the  $\rho$ -competitive subroutine  $A_\rho(2m(t_i))$  is guaranteed to find a feasible solution given a factor  $\rho$  times more machines than optimal, i.e.,  $2\rho m(t_i)$ .

It remains to compare the number of machines opened by **Double** with the optimal number of machines  $m$ , which is at least  $m(t_k)$ . By construction it holds that  $2m(t_i) \leq m(t_{i+1})$ , which implies by recursive application that

$$2m(t_i) \leq \frac{m(t_k)}{2^{k-i-1}}. \quad (3)$$

The total number of machines opened by **Double** is

$$\sum_{i=0}^k 2\rho m(t_i) \leq \sum_{i=0}^k \rho \frac{1}{2^{k-i-1}} m(t_k) = 2\rho \sum_{i=0}^k \frac{1}{2^i} m(t_k) \leq 4\rho m,$$

which concludes the proof.  $\square$

### A.3 Proof of Theorem 3

Recall that EDF is the algorithm that schedules  $m' = \rho m$  unfinished jobs with the smallest deadline at any time  $t$ . Specifically, if there are multiple jobs with the same deadline, we assume that EDF will break the tie arbitrarily. Notice that the release dates, deadlines, and processing times of all jobs are integers. Thus, if EDF decides to process a job at time  $t$ , it actually processes the job during  $[t, t+1)$ , i.e., EDF does not change its decision at fractional time points. Therefore, throughout this subsection, we only consider algorithms that make decisions at integral time points and denote by **A** any such algorithm.

We give some notations. We denote by  $w_A(t)$  the total processing volume of jobs that **A** assigns to  $[t, t+1)$  (or simply to  $t$ ). Specifically, if algorithm **A** makes decisions only at integral time points, this is simply the number of assigned jobs. Further, we let  $p_j^A(t)$  denote the remaining processing time of job  $j$  at time  $t$ ,  $\ell_j^A(t) = d_j - t - p_j^A(t)$ . If it is clear from context, we may drop the superscript **A** and simply write  $p_j(t)$  and  $\ell_j(t)$ . We let  $W_A(t)$  denote the total remaining processing time of all unfinished jobs (including those not yet released) at time  $t$ , i.e.,  $W_A(t) = \sum_{j \in J} p_j - \sum_{s=0}^{t-1} w_A(s)$ . We use **OPT** to indicate an optimal schedule.

A job is called active if it is released but not finished. An algorithm for the (semi-)online machine minimization problem is called *busy* if at all times  $t$ ,  $w_A(t) < m'$  implies that at time  $t$  there are exactly  $w_A(t)$  active jobs, where  $m'$  is the number of machines used by the algorithm.

We prove Theorem 3 by establishing a contradiction based on the workload that EDF when missing deadlines must assign to some interval. We do so by using the following workload inequality, which holds for arbitrary busy algorithms. Here,  $d_{\max}$  denotes the maximum deadline that occurs in the considered instance.

**Lemma 9.** *Let every job be  $\alpha$ -loose,  $\rho \geq 1/(1-\alpha)^2$  and let **A** be a busy algorithm for semi-online machine minimization using  $\rho m$  machines. Assume that  $\ell_j^A(t') \geq 0$  holds, for all  $t' \leq t$  and  $j$ . Then*

$$W_A(t) \leq W_{\text{OPT}}(t) + \frac{\alpha}{1-\alpha} \cdot m \cdot (d_{\max} - t).$$

*Proof.* We prove by induction. The base is clear since  $W_A(0) = W_{\text{OPT}}(0)$ . Now assume the lemma holds for all  $t' < t$ . There are three possibilities.

- Case 1: We have  $w_A(t) \leq \alpha/(1-\alpha) \cdot m$ . By the fact that A is busy, there is an empty machine at time  $t$ , meaning that there are at most  $\alpha/(1-\alpha) \cdot m$  active jobs. Given that  $\ell_j^A(t') \geq 0$  for all  $t' \leq t$ , each of the active jobs has a remaining processing time of no more than  $d_{\max} - t$ , implying that  $W_A(t)$  is bounded by  $\alpha/(1-\alpha) \cdot m \cdot (d_{\max} - t)$  plus the total processing time of the jobs that are not released. As the latter volume cannot exceed  $W_{\text{OPT}}(t)$ , the claim follows.
- Case 2: We have  $w_A(t) \geq 1/(1-\alpha) \cdot m$ . According to the induction hypothesis, it holds that  $W_A(t-1) \leq W_{\text{OPT}}(t-1) + \alpha/(1-\alpha) \cdot m \cdot (d_{\max} - t + 1)$ . Using  $w_{\text{OPT}}(t) \leq m$ , we get  $W_{\text{OPT}}(t) \geq W_{\text{OPT}}(t-1) - m$ . For Algorithm A, it holds that  $W_A(t) = W_A(t-1) - w_A(t-1) \leq W_A(t-1) - 1/(1-\alpha) \cdot m$ . Inserting the first inequality into the second one proves the claim.
- Case 3: We have  $\alpha/(1-\alpha) \cdot m < w_A(t) < 1/(1-\alpha) \cdot m$ . Again by the fact that A is busy, there is an empty machine at time  $t$ , i.e., there are less than  $1/(1-\alpha) \cdot m$  active jobs. Again distinguish two cases.
- Case 3a: We have  $p_j(t) \leq \alpha(d_j - t)$  for all active jobs. Then the total remaining processing time of active jobs is bounded by  $\alpha(d_{\max} - t) \cdot 1/(1-\alpha) \cdot m$ . Plugging in the total processing time of unreleased jobs, which is bounded by  $W_{\text{OPT}}(t)$ , we again get the claim.
- Case 3b: There exists an active job  $j$  with  $p_j(t) > \alpha(d_j - t)$ . Using that  $j$  is  $\alpha$ -loose, i.e.,  $p_j \leq \alpha(d_j - r_j)$ , we get that  $j$  is not processed for at least  $(1-\alpha)(t - r_j)$  time units in the interval  $[r_j, t)$ . As Algorithm A is busy, this means that all machines are occupied at these times, yielding

$$\begin{aligned}
W_A(t) &\leq W_A(r_j) - (1-\alpha)(t - r_j) \cdot \rho m \\
&\leq W_{\text{OPT}}(r_j) + \frac{\alpha}{1-\alpha} \cdot m \cdot (d_{\max} - r_j) - (1-\alpha)(t - r_j) \cdot \rho m \\
&\leq W_{\text{OPT}}(r_j) + \frac{\alpha}{1-\alpha} \cdot m \cdot (d_{\max} - t) - m \cdot (t - r_j),
\end{aligned}$$

where the second inequality follows by the induction hypothesis for  $t = r_j$ , and the third one follows by  $\rho \geq 1/(1-\alpha)^2$ . Lastly, the feasibility of the optimal schedule implies

$$W_{\text{OPT}}(t) \geq W_{\text{OPT}}(r_j) - m \cdot (t - r_j),$$

which in turn implies the claim by plugging it into the former inequality.  $\square$

We now apply this lemma in the proof of the theorem.

*Proof (Theorem 3).* Let  $\rho m$  be the number of machines EDF uses and consider the schedule produced by EDF from an arbitrary instance  $J$  only consisting of  $\alpha$ -loose jobs. We have to prove that every job is finished before its deadline if  $\rho = 1/(1-\alpha)^2$ . To this end, suppose that EDF fails. Among the jobs that are missed, let  $j^*$  be one of those with the earliest release date.

First observe that we can transform  $J$  into  $J' \subseteq J$  such that  $j^* \in J'$ ,  $\max_{j \in J'} d_j = d_{j^*}$  and EDF still fails on  $J'$ . For this purpose, simply define  $J' = \{j \in J \mid d_j \leq d_{j^*}\}$  and notice that we have only removed jobs with a later deadline than  $\max_{j \in J'} d_j$ . This implies that every job from  $J'$  receives the same processing in the EDF schedule of  $J'$  as in the one of  $J$ .

We can hence consider  $J'$  instead of  $J$  from now on. In the following, we establish a contradiction on the workload during the time interval  $[r_{j^*}, d_{j^*})$ . The first step towards this

is applying Lemma 9 for an upper bound. Also making use of the feasibility of the optimal schedule, we get:

$$\begin{aligned} W_{\text{EDF}}(r_{j^*}) &\leq W_{\text{OPT}}(r_{j^*}) + \frac{\alpha}{1-\alpha} \cdot m \cdot (d_{\max} - r_{j^*}) \\ &\leq \frac{1}{1-\alpha} \cdot m \cdot (d_{\max} - r_{j^*}). \end{aligned}$$

We can, however, also lower bound this workload. Thereto, note that job  $j^*$  is not processed for at least  $(1-\alpha)(d_{j^*} - r_{j^*})$  time units where  $d_{j^*} = d_{\max}$ , implying that all machines must be occupied by then, i.e.,

$$W_{\text{EDF}}(r_{j^*}) > (1-\alpha) \cdot \rho m \cdot (d_{\max} - r_{j^*}).$$

If we compare the right-hand sides of the two inequalities, we arrive at a contradiction if  $\rho \geq 1/(1-\alpha)^2$ , which concludes the proof.  $\square$

## B Missing Material from Section 5

### B.1 Analysis of the Job Assignment

We show that the algorithm does not fail to assign any job to a group for some  $\tau = \mathcal{O}(m \log m)$ . The basic idea is the same as that for laminar instances, that is, we again select out a critical job set if the algorithm fails to assign a job, and we derive a contradiction by a load argument from this set.

In particular, for a fixed job  $j^*$  whose assignment fails, we iteratively define its *failure set*  $F \subseteq J$ , which is the union of the  $\tau + 1$  different sets  $F_0, \dots, F_\tau$ . To define  $F_0$ , we also use auxiliary sets  $F_0^0, \dots, F_0^\tau$ .

We initialize  $F_0^\tau := \{j^*\}$ . Given  $F_0^i$ , we construct  $F_i$  and  $F_0^{i-1}$  in the following way. First,  $F_i$  is defined to be the set of all  $\prec$ -maximal  $i$ -th candidates of jobs in  $F_0^i$ . Subsequently,  $F_0^{i-1}$  is constructed by adding all the  $i$ -th users of jobs in  $F_i$  to  $F_0^i$ . Formally,

$$\begin{aligned} F_i &:= M_{\prec}(\{c_i(j) \mid j \in F_0^i\}) \\ \text{and } F_0^{i-1} &:= F_0^i \cup \bigcup_{j \in F_i} U_i(j), \end{aligned}$$

where  $M_{\prec}$  is the operator that picks out the  $\prec$ -maximal elements from a set of jobs:  $M_{\prec}(S) := \{j \in S \mid \nexists j' : j \prec j'\}$ . After  $\tau$  such iterations, that is, when  $F_0^0, F_1, \dots, F_\tau$  have been computed, we set  $F_0 := M_{\prec}(F_0^0)$ .

Again, we have the following lemma whose proof is exactly the same as that of Lemma 3. Notice that, due to the difference in the bin sizes, here  $F_0$   $\gamma$ -blocks each  $F_i$  where  $\gamma = 1/(qm\tau)$ , which is different from Lemma 3.

**Lemma 10.** *Every failure set  $F$  has the following properties:*

- (i) *We have  $F_0 \prec \dots \prec F_\tau$ .*
- (ii) *The sets  $F_0, \dots, F_\tau$  are pairwise disjoint.*
- (iii) *For all  $i = 1, \dots, \tau$ ,  $F_0$   $\gamma$ -blocks  $F_i$  where  $\gamma = 1/(qm\tau)$ .*

The following lemma is also true via the same proof as Lemma 4.



**Lemma 11.** *For any fixed  $\alpha \in (0, 1)$ , there exists  $k = \mathcal{O}(m)$  with the following property. If  $F_i$   $\gamma$ -blocks  $F_h$ , then  $F_{i+k}$   $2\gamma$ -blocks  $F_h$ , for all  $i + k < h$ .*

**Lemma 12.** *There exists  $g = \mathcal{O}(m^2 \log m)$  such that the job assignment never fails.*

*Proof.* We will choose  $\tau = \mathcal{O}(m \log m)$  such that we get a contradiction to the existence of the failure set  $F$ . Using Lemma 10 (iii) and Lemma 11, we can choose  $\lambda = \mathcal{O}(m \log(qm\tau))$  such that  $F_\lambda$  2-blocks every  $F_i$  with  $i > \lambda$ . Now we lower bound the total contribution of these  $F_i$  to  $I(F_\lambda)$ : For some  $j \in F_i$  where  $i \geq \lambda$ , its contribution to  $I(F_\lambda)$  is

$$|I(F_\lambda) \cap I(j)| - \ell_j \geq \frac{|I(F_\lambda) \cap I(j)|}{2}$$

since  $|I(F_\lambda) \cap I(j)| \geq 2\ell_j$ . Thus we get as total contribution of  $F_i$  to  $F_\lambda$  at least

$$\begin{aligned} \sum_{j \in F_i} \frac{|I(F_\lambda) \cap I(j)|}{2} &\geq \frac{1}{2} \cdot \left| I(F_\lambda) \cap \left( \bigcup_{j \in F_k} I(j) \right) \right| \\ &= \frac{|I(F_\lambda) \cap I(F_i)|}{2} \geq \frac{|I(F_\lambda)|}{2}, \end{aligned}$$

where the last inequality follows from  $F_\lambda \prec F_i$ . Therefore, the total contribution of all  $F_i$  for  $i > \lambda$  is at least  $(\tau - \lambda)|I(F_\lambda)|/2$ , which is at most  $m|I(F_\lambda)|$  by Theorem 1. Thus, we must have  $\tau \leq \lambda + 2m = \mathcal{O}(m \log(qm\tau))$ , and appropriately choosing  $\tau = \mathcal{O}(m \log m)$  yields a contradiction. By Lemma 8, we can choose  $g = \mathcal{O}(m\tau) = \mathcal{O}(m^2 \log m)$ .  $\square$

We now obtain Theorem 6 by putting Lemmas 7 and 12 together.